ABSTRACT

A longest matching prefix is identified based on a search of ranges corresponding to the prefixes. These prefixes may correspond to any values, such as, but not limited to Internet Protocol (IP) or other network addresses, character strings, etc. A first-level tiny tree is traversed to identify a next-level tiny tree. A root node of the next-level tiny tree is retrieved. The root node typically includes a back value to identify a matching prefix should no matching prefix be identified within said particular tree. The particular next-level tiny tree is traversed to either identify a matching prefix or a no match condition, with the traversing said particular next-level tiny tree typically including comparing the lookup value with one or more of the plurality of keys. The prefix identified based on the back value is identified as the longest matching prefix if said traversing resulted in said no match condition else the matching prefix.